Master Project

Model Learning in Active Inference: An Evaluation of Algorithms and their Behaviour

Laurens Y. D. Žnidaršič, MSc S&C and MSc ME-BMD, TU Delft
L.Y.D.Znidarsic@student.tudelft.nl
Martijn Wisse, Department of BioMechaninical Engineering, TU Delft
M.Wisse@tudelft.nl
Peyman Mohajerin Esfahani, Delft Center for Systems and Control, TU Delft
P.MohajerinEsfahani@tudelft.nl

Context

For decades, neuroscientists have been actively theorizing how one of the most complex systems we know, the human brain, works and interacts with the body and the world. Major breakthroughs have led to the understanding that our brain relies heavily on filtering and predicting the state of the world around us, i.e. perception and interpretation.

Active Inference

One relatively novel theory, which has grown to be the closest one we have towards a unifying brain theory, is that of Active Inference. It provides a fundamentally new level of integration between perception and action; where classical system- and control point of view sees a filter and a controller as separate entities, active merges them via Free-Energy minimization (FEM).

Even though the figure on the right, which schematically depicts the Active Inference framework, contains separate blocks for a filter and a controller, they are in fact minimizing the same cost function, and they should therefore be seen as a unified entity.

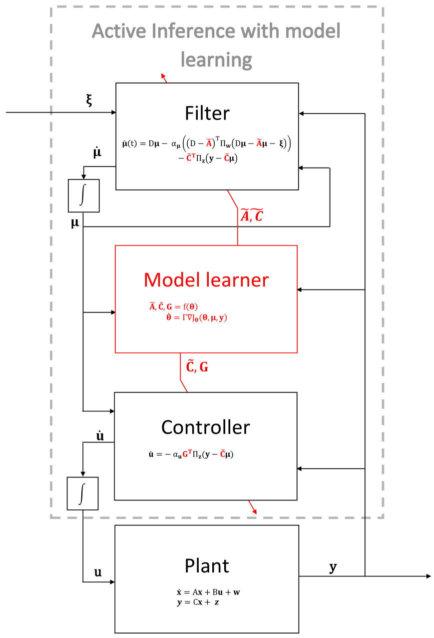

Figure 1: Schematic of the Active Inference Framework. $\boldsymbol{\xi}$ relates to a reference and $\boldsymbol{\mu}$ to a filtered state. $\boldsymbol{u}$ and $\boldsymbol{y}$ are the plant in- and outputs respectively. The filter and controller are driven by gradient descent on a cost function called the Free energy, which consists of the terms sensory surprise and internal consistency. The model learner updates the agent's belief about it's own behaviour (the model) based on the (filtered) state and system input.

Model Learning
The figure also shows, apart from the plant itself, a fourth block: the Model Learner. This is where I am focusing on for my thesis project; the Active Inference framework requires some internal belief about the system, i.e. a forward model. Up until now, this model had always been considered static and known a priori, but as the founding father of Active Inference, Karl Friston, suggests, it does not necessarily have to be.

This model can be learned from the behaviour of the Plant by collecting the input- and output data. The learning algorithm can be either gradient-based or data-driven, it can converge to a final model or be ever-updating, and it can be linear or non-linear, possibly even hierarchical.

Figure 2: The Jackal skid-steering robot

Project tasks

Because of the large number of options, a first step in this research direction will be to list the algorithms available and evaluate their influence, behaviour and performance within the Active Inference Framework. This evaluation will be done on both a theoretical level, e.g. prove stability, convergence and robustness, and on an experimental level, e.g. implement the most promising algorithms. Relating the experimental results of the different algorithms to the way they are constructed will be one of the hardest tasks of this part of the research.

Later on, the most promising algorithms will also be used in state-prediction and control of a skid-steering robot (figure 2). This system is known to contain a high level of Non-Markovian Noise, which is a setting in which Active Inference can theoretically outperform classical control methods. From these experiments the overall performance of the Active Inference Framework with Model Learning will be measured against other control methods.

1. Establish criteria and metrics to evaluate algorithm performance
2. Theorize what influence learning in general has on the overall behaviour of the active inference framework and establish conditions for suitable algorithms
3. Distill algorithms suitable for using withing the Active Inference framework
4. Theorize the influence different algorithms have on the overall behaviour of the active inference framework
5. Experiment in simulation Active inference with model learning on simple linear state-space systems
6. Experiment in simulation Active inference with model learning on simple nonlinear state-space systems
7. Experiment Active inference with model learning on Jackal Robot
8. Measure performance against other (advanced) control methods
9. Relate experimental results to theory